


The Idea in a Nutshell



- Classify a new item based on the classes of its k nearest neighbors.
- Sensitivity to
 - k
 - distance metric
 - how the k points contribute to the decision

Computational Challenges

- Distance metric needs to be tuned (e.g., parameter scaling).
- Given the neighbors, making a decision is tricky.
 - (Mixed) optimization.
- Over-fitting should be avoided.
- Finding the nearest neighbors might be time consuming.
 - Data management/data decomposition/ parallel processing.

Pros

- Simple
- Flexible
- Incremental
- Amenable to parallelization.

Cons

- Too simple
- Sensitive to the distance metric, k , and voting.
- Classification varies with the order data is provided.